

ATTY. DOCKET NO. SAMS01-00188

PATENT

DISTRIBUTED IDENTITY SERVER FOR USE IN  
A TELECOMMUNICATION SWITCH

Inventor(s) :

David Allan Collins  
3301 Bending Oaks Trail  
Garland  
Texas 75044  
Dallas County  
United States citizen

Assignee:

SAMSUNG ELECTRONICS Co., LTD.  
416, Maetan-dong, Paldal-gu  
Suwon-city, Kyungki-do  
Republic of Korea

CERTIFICATE OF EXPRESS MAIL

I hereby certify that this correspondence, including the attachments listed, is being mailed in an envelope addressed to Commissioner of Patents and Trademarks, Washington, DC 20231, using the Express Mail Post Office to Addressee service of the United States Postal Service on the date shown below.

Kathy Longenecker

Printed Name of Person Mailing

*Kathy Longenecker*  
Signature of Person Mailing

ET838009796US

Express Mail Receipt No.

*12/31/01*  
Date

William A. Munck  
John T. Mockler  
Novakov Davis & Munck, P.C.  
Three Galleria Tower  
13155 Noel Road, Suite 900  
Dallas, Texas 75240  
(214) 922-9221

DISTRIBUTED IDENTITY SERVER FOR USE IN  
A TELECOMMUNICATION SWITCH

CROSS-REFERENCE TO RELATED APPLICATIONS

The present invention is related to those disclosed in the  
5 following United States Non-Provisional Patent Applications:

- 1) [Docket No. SAMS01-00186] filed concurrently herewith,  
entitled "SYSTEM AND METHOD FOR DISTRIBUTED CALL PROCESSING  
USING LOAD SHARING GROUPS";
- 10 2) [Docket No. SAMS01-00187] filed concurrently herewith,  
entitled "SYSTEM AND METHOD FOR DISTRIBUTED CALL PROCESSING  
USING A DISTRIBUTED TRUNK IDLE LIST";
- 3) [Docket No. SAMS01-00189] filed concurrently herewith,  
entitled "SYSTEM AND METHOD FOR PROVIDING A SUBSCRIBER  
DATABASE USING GROUP SERVICES IN A TELECOMMUNICATION SYSTEM."

15 The above applications are commonly assigned to the assignee  
of the present invention. The disclosures of these related patent  
applications are hereby incorporated by reference for all purposes  
as if fully set forth herein.

**TECHNICAL FIELD OF THE INVENTION**

The present invention is directed, in general, to telecommunication systems and, more specifically, to a telecommunication system that uses group services to implement a distributed identity server.

**BACKGROUND OF THE INVENTION**

Wireless service providers continually try to create new markets and to expand existing markets for wireless services and equipment. One important way to accomplish this is to improve the performance of wireless network equipment while making the equipment cheaper and more reliable. Doing this allows wireless service providers to reduce infrastructure and operating costs while maintaining or even increasing the capacity of their wireless networks. At the same time, the service providers are attempting to improve the quality of wireless service and increase the quantity of services available to the end-user.

The mobile switching of a wireless network provides connections between a number of wireless network base stations and the public switched telephone network. Calls originated by or terminated at a cell phone or other mobile station are handled in the mobile station by a number of call processing client

applications. A conventional mobile station typically contains a large switching fabric controlled by a main processing unit (MPU) that contains a large number of data processors and associated memories, often in the form of ASIC chips. Each of these MPU  
5 processors contains a call process client application for controlling the flow of control signals of a single call. Each call process client application in turn communicates with a call process server application that controls the flow of control signals for a large number of calls.

10 Thus, when a particular event occurs during a phone call (e.g., the call set-up, the invocation of three-way calling, call disconnection, or the like), control signals associated with the event are relayed from the mobile station to the call process client application in the mobile switching center (MSC). This call  
15 processing client application then relays the control signals to the call process server application, which actually performs the call processing service requested by the control signals.

Unfortunately, in large capacity systems, bottlenecks may develop around the call process server applications. Each call  
20 process client application must communicate with a particular piece of server hardware that is executing the call process server application. Due to the random nature of the starts, stops and

other events in phone calls, in a large system, some servers may be near capacity and develop bottlenecks, while other servers still have plenty of adequate bandwidth. Moreover, a system failure in a particular piece of server hardware results in the loss of all call processes being handled by a call process server application being executed on the failed server.

When a new call begins (technical term is "originated"), there is a need to uniquely identify the call during its entire lifetime. The unique identity is referred to as the "call identity" and is a unique number. If any events occur during the lifetime of the call (such as answer trigger, various timeouts, hang-ups, and the like), the event is delivered to the call processing application that is handling the specific call. The call identity value is used to locate and forward the events to the correct call processing application. Conventional systems use a single server to be responsible for allocating a call identity value. The prior art server allocates any number when requested and marks that number as being in use. When the call ends, the call identity number is released and is marked by the server as being unallocated. That number may then be reused for another call. For higher availability, the server would have a backup so that if the server failed, the backup could take over. However, this approach

suffers from being a bottle neck and having poor scalability.

Therefore, there is a need for improved wireless network equipment and services. In particular, there is a need for mobile switching centers that are highly reliable and minimally susceptible to bottleneck conditions during periods of high call traffic volume. More particularly, there is a need for an improved call identity server architecture for use in mobile switching centers and other similar switching devices.

## SUMMARY OF THE INVENTION

To address the above-discussed deficiencies of the prior art, it is a primary object of the present invention to provide a controller for allocating call identity values to call connections associated with a switch, wherein the switch handles call connections between calling devices and called devices on trunk lines associated with the switch. According to an advantageous embodiment of the present invention, the controller comprises N call application nodes for executing a plurality of identity server applications that allocate call identity values to the call connections, wherein a first one of the plurality of identity server applications is executed on a first one of the N call application nodes and is associated with a second one of the identity server applications executed on a second one of the N call application nodes separate from the first call application node, the first and second identity server applications forming a load sharing group server application, and wherein the load sharing group server application receives a call identity request from a new call process being executed in the switch and selects one of the first and second identity server applications to allocate a call identity value to a new call connection associated with the call identity request according to a load distribution algorithm.

According to one embodiment of the present invention, the first identity server application allocates call identity values having a first contiguous range and the second identity server application allocates call identity values having a second  
5 contiguous range different than the first contiguous range.

According to another embodiment of the present invention, the load distribution algorithm distributes new call identity requests in an alternating manner between the first and second identity server applications.

10 According to still another embodiment of the present invention, the load distribution algorithm distributes new call identity requests according to a current processing load of the first identity server application and a current processing load of the second identity server application.

15 According to yet another embodiment of the present invention, the load distribution algorithm distributes the new call identity requests in order to maintain the current processing load of the first identity server application at a level substantially equal to the current processing load of the second identity server  
20 application.

According to a further embodiment of the present invention, the first identity server application comprises a first primary-



5 backup group server application, wherein the first primary-backup group server application comprises a first primary identity server application executed on the first call application node and a first backup identity server application associated with the first primary identity server application.

10 According to a still further embodiment of the present invention, call state information associated with the first primary identity server application is mirrored to the first backup identity server application associated with the first primary identity server application.

15 According to a yet further embodiment of the present invention, the first backup identity server application resides on the first call application node.

20 In one embodiment of the present invention, the first backup identity server application resides on a call application node separate from the first call application node.

In another embodiment of the present invention, the second identity server application comprises a second primary-backup group server application, wherein the second primary-backup group server application comprises a second primary identity server application executed on the second call application node and a second backup identity server application associated with the second primary

identity server application.

In still another embodiment of the present invention, state information associated with the second primary call process is mirrored to the second backup call process associated with the  
5 second primary call process.

in yet another embodiment of the present invention, the second backup identity server application resides on the second call application node.

In a further embodiment of the present invention, the second  
10 backup identity server application resides on a call application node separate from the second call application node.

The foregoing has outlined rather broadly the features and technical advantages of the present invention so that those skilled in the art may better understand the detailed description of the  
15 invention that follows. Additional features and advantages of the invention will be described hereinafter that form the subject of the claims of the invention. Those skilled in the art should appreciate that they may readily use the conception and the specific embodiment disclosed as a basis for modifying or designing  
20 other structures for carrying out the same purposes of the present invention. Those skilled in the art should also realize that such equivalent constructions do not depart from the spirit and scope of

the invention in its broadest form.

Before undertaking the DETAILED DESCRIPTION OF THE INVENTION below, it may be advantageous to set forth definitions of certain words and phrases used throughout this patent document: the terms

5 "include" and "comprise," as well as derivatives thereof, mean inclusion without limitation; the term "or," is inclusive, meaning and/or; the phrases "associated with" and "associated therewith," as well as derivatives thereof, may mean to include, be included within, interconnect with, contain, be contained within, connect to

10 or with, couple to or with, be communicable with, cooperate with, interleave, juxtapose, be proximate to, be bound to or with, have, have a property of, or the like; and the term "controller" means any device, system or part thereof that controls at least one operation, such a device may be implemented in hardware, firmware

15 or software, or some combination of at least two of the same. It should be noted that the functionality associated with any particular controller may be centralized or distributed, whether locally or remotely. Definitions for certain words and phrases are provided throughout this patent document, those of ordinary skill

20 in the art should understand that in many, if not most instances, such definitions apply to prior, as well as future uses of such defined words and phrases.

## BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention, and the advantages thereof, reference is now made to the following descriptions taken in conjunction with the accompanying drawings, wherein like numbers designate like objects, and in which:

FIGURE 1 illustrates an exemplary wireless network according to one embodiment of the present invention;

FIGURE 2 illustrates an exemplary mobile switching center in greater detail according to one embodiment of the present invention;

FIGURE 3 illustrates selected portions of a mobile switching center that perform distributed call processing using group services according to the principles of the present invention;

FIGURE 4A is an illustration of server side internal group policy classes according to an exemplary embodiment of the present invention;

FIGURE 4B is an illustration of a client side internal client policy architecture according to an exemplary embodiment of the present invention; and

FIGURE 4C is an illustration of a load sharing client side policy internal architecture according to an exemplary embodiment of the present invention.

## DETAILED DESCRIPTION OF THE INVENTION

FIGURES 1 through 4, discussed below, and the various embodiments used to describe the principles of the present invention in this patent document are by way of illustration only and should not be construed in any way to limit the scope of the invention. Those skilled in the art will understand that the principles of the present invention may be implemented in any suitably arranged telecommunications network.

In the disclosure that follows, a group services framework for performing various distributed call processing functions is implemented in a mobile switching center of a wireless communication network. This is by way of illustration only and should not be construed so as to limit the scope of the invention. Those skilled in the art will understand that the group services framework described below may be implemented in other types of telecommunication devices, including many varieties of switches, routers and the like.

FIGURE 1 illustrates exemplary wireless network 100 according to one embodiment of the present invention. Wireless network 100 comprises a plurality of cell sites 121-123, each containing one of the base stations, BS 101, BS 102, or BS 103. Base stations 101-103 communicate with a plurality of mobile stations (MS) 111-114

over, for example, code division multiple access (CDMA) channels. Mobile stations 111-114 may be any suitable wireless devices, including conventional cellular radiotelephones, PCS handset devices, personal digital assistants, portable computers, or metering devices. The present invention is not limited to mobile devices. Other types of access terminals, including fixed wireless terminals, may be used. However, for the sake of simplicity, only mobile stations are shown and discussed hereafter.

Dotted lines show the approximate boundaries of the cell sites 121-123 in which base stations 101-103 are located. The cell sites are shown approximately circular for the purposes of illustration and explanation only. It should be clearly understood that the cell sites may have other irregular shapes, depending on the cell configuration selected and natural and man-made obstructions.

As is well known in the art, cell sites 121-123 are comprised of a plurality of sectors (not shown), each sector being illuminated by a directional antenna coupled to the base station. The embodiment of FIGURE 1 illustrates the base station in the center of the cell. Alternate embodiments position the directional antennas in corners of the sectors. The system of the present invention is not limited to any one cell site configuration.

In one embodiment of the present invention, BS 101, BS 102, and BS 103 comprise a base station controller (BSC) and one or more base transceiver subsystem(s) (BTS). Base station controllers and base transceiver subsystems are well known to those skilled in the art. A base station controller is a device that manages wireless communications resources, including the base transceiver stations, for specified cells within a wireless communications network. A base transceiver subsystem comprises the RF transceivers, antennas, and other electrical equipment located in each cell site. This equipment may include air conditioning units, heating units, electrical supplies, telephone line interfaces, and RF transmitters and RF receivers. For the purpose of simplicity and clarity in explaining the operation of the present invention, the base transceiver subsystem in each of cells 121, 122, and 123 and the base station controller associated with each base transceiver subsystem are collectively represented by BS 101, BS 102 and BS 103, respectively.

BS 101, BS 102 and BS 103 transfer voice and data signals between each other and the public switched telephone network (PSTN) (not shown) via communication trunk lines 131, mobile switching center (MSC) 140, and communication trunk lines 132. Trunk lines 131 also provide connection paths to transfer control signals

between MSC 140 and BS 101, BS 102 and BS 103 that are used to establish connections for voice and data circuits between MSC 140 and BS 101, BS 102 and BS 103 over communication trunk lines 131 and between MSC 140 and the Internet or the PSTN over communication trunk lines 132. In some embodiments of the present invention, communication trunk lines 131 may be several different data links, where each data link couples one of BS 101, BS 102, or BS 103 to MSC 140.

Trunk lines 131 and 132 comprise one or more of any suitable connection means, including a T1 line, a T3 line, a fiber optic link, a network packet data backbone connection, or any other type of data connection. Those skilled in the art will recognize that the connections on trunk lines 131 and 132 may provide a transmission path for transmission of analog voice band signals, a digital path for transmission of voice signals in the pulse code modulated (PCM) format, a digital path for transmission of voice signals in an Internet Protocol (IP) format, a digital path for transmission of voice signals in an asynchronous transfer mode (ATM) format, or other suitable connection transmission protocol. Those skilled in the art will recognize that the connections on trunk lines 131 and 132 may provide a transmission path for transmission of analog or digital control signals in a suitable



signaling protocol.

FIGURE 2 illustrates exemplary mobile switching center 140 in greater detail according to one embodiment of the present invention. MSC 140 includes interconnecting network 200, among  
5 other things. Interconnecting network 200 comprises switch fabric 205 and switch controller 210, which together provide switch paths between communication circuits in trunk lines 131 and 132. MSC 140 provides services and coordination between the subscribers in wireless network 100 and external networks, such as the PSTN or  
10 Internet. Mobile switching centers similar to MSC 140 are well known to those skilled in the art.

When a wireless network subscriber turns on his or her mobile station (e.g., cell phone) or fixed access terminal, radio messages over the air interface inform the base station that the mobile  
15 station (or fixed access terminal) is joining the network. However, a connection is not automatically made to voice or data traffic carrying circuits in trunk lines 131-132. A voice or data traffic connection to the public switched telephone network (PSTN) or the Internet is not needed until the subscriber places a call  
20 (e.g., dials a phone number) or accesses the Internet.

However, even when the phone is idle, certain information about the subscriber (i.e., subscriber data) must be retrieved and

stored in either the base station or in MSC 140, or both, in order to authenticate the subscriber, gather billing information, identify the services available to the subscriber, determine capabilities of the mobile station, and the like. The control signals (as opposed to voice and data traffic) required to do this are also carried over trunk lines 131 and 132. After the subscriber data is stored in memory in MSC 140, it is available for use by a variety of call processing client (CPC) applications that may be initiated by the subscriber or another device while the mobile station is still active.

For example, when MS 111 is first turned ON, a call process is set up in MSC 140 for MS 111 and subscriber data (e.g., billing information) is stored in MSC 140 that may be accessed by the call process or other call applications that provide particular types of call services. If the subscriber dials a phone number on MS 111 or a call is received from the PSTN directed to MS 111, the call process for MS 111 handles the establishment of a call connection on one of the trunk lines in trunk line 131 and one of the trunk lines in trunk line 132. The MS 111 call process executed in MSC 140 maintains all state information related to the call and to MS 111 and handles all other applications required by MS 111, including three-way calls, voice mail, call disconnection, and the

like.

In order to handle a large amount of call traffic, it is necessary to distribute the many active call processes and call service applications handled by MSC 111 across a number of call application nodes. The call services may include application for  
5 accessing a subscriber database, selecting (or de-selecting) trunk, lines, retrieving and maintaining call identity information, and the like. The present invention provides methods and apparatuses for distributing call processes and call service applications  
10 across multiple call application nodes in a highly reliable and redundant manner. This is accomplished by a distributed network of redundant servers in which call traffic is distributed in order to increase the call-handling capacity of MSC 140. The redundancy of the distributed servers is transparent to both the call process  
15 client applications that require a service and the call process server applications that provide the service. It also decreases the complexity of both the client and server applications.

FIGURE 3 illustrates in greater detail selected portions of exemplary mobile switching center 140 that perform distributed call processing using group services in accordance with the principles  
20 of the present invention. MSC 140 comprises main processing unit (MPU) 310, system manager node 1 (SYSMGR1), optional system manager

node 2 (SYSMGR2), and master database 320. MSC 140 also comprises a plurality of call application nodes (CANs), including CAN1, CAN2, and CAN3, and a plurality of local storage devices (SDs), namely SD1, SD2, and SD3, that are associated with CAN1, CAN2 and CAN3.

5 Master database 320 may be used as a master software repository to store databases, software images, server statistics, log-in data, and the like. SD1-SD3 may be used to store local capsules, transient data, and the like.

Each one of system manager nodes 1 and 2 and CAN1-CAN3  
10 executes a configuration management (CM) process that sets up each node with the appropriate software and configuration data upon initial start-up or after a reboot. Each node also executes a node monitor (NM) process that loads software and tracks processes to determine if any process has failed. System manager nodes 1 and 2  
15 execute a first arbitrary process, P1, and system manager node 1 also executes a second arbitrary process, P2.

In accordance with the principles of the present invention, call application nodes 1-3 (CAN1-CAN3) also execute a number of call process (CP) server applications organized as primary and  
20 backup processes that are available as distributed group services to 1 to N call process client (CPC) applications, namely CPC APP1 - CPC APPn in main processing unit 310. The N call application nodes

(e.g., CAN1-CAN3) are separate computing nodes comprising a processor and memory that provide scalability and redundancy by the simple addition of more call application nodes.

Each of the N call process client (CPC) applications, namely  
5 CPC APP1 - CPC APPn in MPU 310 handles the control signals and messages related to a single call associated with a mobile station. Each of CPC APP1 - CPC APPn establishes a session with a load sharing group, which assigns the call to a particular one of the primary-backup group call process server applications, CP1, CP2, or  
10 CP3. The selected call process server application actually performs the call process services/functions requested by the call process client application.

In the illustrated embodiment, three exemplary call process server applications are being executed, namely CP1, CP2, and CP3.  
15 Each of these processes exists as a primary-backup group. Thus, CP1 exists as a primary process, CP1(P), and a backup process, CP1(B). Similarly, CP2 exists as a primary process, CP2(P), and a backup process, CP2(B), and CP3 exists as a primary process, CP3(P), and a backup process, CP3(B). In the illustrated  
20 embodiment, CP1(P) and CP1(B) reside on different call application nodes (i.e., CAN1 and CAN2). This is not a strict requirement: CP1(P) and CP1(B) may reside on the same call application node

(e.g., CAN1) and still provide reliability and redundancy for software failures of the primary process, CP1(P). However, in a preferred embodiment of the present invention, the primary process and the backup process reside on different call application nodes, thereby providing hardware redundancy as well as software redundancy. Thus, CP1(P) and CP1(B) reside on CAN1 and CAN2, CP2(P) and CP2(B) reside on CAN2 and CAN3, and CP3(P) and CP3(B) reside on CAN3 and CAN1.

Together, CP1, CP2 and CP3 form a supergroup for load sharing purposes. Thus, CP1(P) and CP1(B), CP2(P) and CP2(B), and CP3(P) and CP3(B) are part of a first load sharing group (LSG1), indicated by the dotted line boundary. Additionally, CAN1-CAN3 host three other load sharing groups, namely, LSG2, LSG3, and LSG4. LSG2 comprises two trunk idle list (TIL) server applications, namely TIL1 and TIL2. TIL1 exists as a primary process, TIL1(P), on CAN2 and a backup process, TIL1(B), on CAN3. TIL2 exists as a primary process, TIL2(P), on CAN3 and a backup process, TIL2(B), on CAN2. Similarly, LSG3 comprises two identity server (IS) applications, namely IS1 and IS2. IS1 exists as a primary process, IS1(P), on CAN1 and a backup process, IS1(B), on CAN2 and IS2 exists as a primary process, IS2(P), on CAN2 and a backup process, IS2(B), on CAN1. Finally, LSG4 comprises two subscriber database (SDB) server

applications, namely SDB1 and SDB2. SDB1 exists as a primary process, SDB1(P), on CAN2 and a backup process, SDB1(B), on CAN3 and SDB2 exists as a primary process, SDB2(P), on CAN3 and a backup process, SDB2(B), on CAN2.

5 A group service provides a framework for organizing a group of distributed software objects in a computing network. Each software object provides a service. In addition, the group service framework provides enhanced behavior for determining group membership, deciding what actions to take in the presence of faults, and controlling unicast, multicast, and groupcast  
10 communications between members and clients for the group. A group utilizes a policy to enhance the behavior of the services provided by the group. Some of these policies include primary-backup for high service availability and load sharing for distributing the  
15 loading of services within a network.

Call processing server applications, such as CP1-CP3, IS1-IS2, and TIL1-TIL2, located within a computing network provide services that are invoked by client applications, such as CPC APP1 - CPC APPn. As shown in FIGURE 3, the call processing server  
20 applications are organized into primary-backup groups configured as a 1+1 type of primary-backup group. There are multiple numbers of these primary-backup groups and the exact number is scalable

according to the number of processes and/or computing nodes (CANS) that are used. All of the primary-backup groups are themselves a member of a single load sharing group (e.g., LSG1, LSG2, LSG3, LSG4).

5 It is important to note that while the call process client applications, CPC APP1-CPC APPn, are clients with respect to the call process server applications, CP1, CP2, and CP3, a server application may be a client with respect to another server application. In particular, the call process server applications  
10 CP1-CP3 may be clients with respect to the trunk idle list server applications, TIL1 and TIL2, the subscriber database server applications, SDB1 and SDB2, and the identity server applications, IS1 and IS2.

A client application establishes an interface to the load  
15 sharing group. When a new call indication is received by the client application, the client application establishes a session with the load sharing group according to a client-side load sharing policy. The initial policy is round-robin (i.e., distribution of new calls in sequential order to each CAN), but other policies may  
20 be used that take into account the actual loading of the different primary-backup groups.

The client application associates the session with the new



call and sends messages associated with the call over the session object. The client application also receives messages from the primary-backup group via the session established with the primary-backup group. Only the primary process (e.g., CP1(P)) of the primary-backup group joins the load sharing group (e.g., LSG1). For a variety of reasons, the application containing the primary may be removed from service. The server application may elect to not accept any new calls by leaving the load sharing group. However, the client applications may still maintain their session with the primary-backup group for existing calls. This action is taken because new call traffic may be lost if the singleton primary also fails. New calls are not distributed to the primary-backup group if it leaves the load sharing group.

If the primary of the primary-backup group that is a member of the load sharing group should fail, the backup member is informed that the primary member has failed (or left) and then assumes the role of primary member. The responsibility for these actions must be performed by the server application. It is the responsibility of the Group Service to inform the backup member that the primary member has failed or left.

As part of an online software upgrade process, one or more applications containing primary-backup groups may be removed from

service, brought down, and then brought back up using a new version of software code. These groups, if their interface has not changed, join the existing load sharing group. When first started, it is required that the client interface be capable of throttling the call traffic to specific primary-backup groups. The traffic throttling is expressed as a percentage varying from 0% (no calls) to 100%. All new calls that would have been scheduled according to the scheduling algorithm are handled by this session. The throttling factor is initialized to 100% for any primary-backup group that joins the load sharing group. During on-line software upgrades, the throttling factor is adjusted to start with the no-calls case for the new software version. Any client application for the load sharing group may establish a session with a specific primary-backup group. The client may then change the throttling factor at any time. When the throttling factor is changed, all client session interfaces receive via multicast the changed throttling factor. As the throttling factor is increased, the call process server applications with the new software version may receive increasing amounts of call traffic.

Call processing communications from the client applications to the call processing server primary-backup groups must support a very high volume of calls. The group software utilizes an internal

transport consisting of a multicasting protocol (simple IP multicast) and optionally a unicasting protocol. The unicasting protocol may be TCP/IP, SCTP, or other transport protocol. The multicast protocol is used for internal member communications relating to membership, state changes, and fault detection. In the absence of unicast transport, the multicast protocol is used for client/server communication streams. The unicast protocol, when provided, is used to provide a high-speed stream between clients and servers. The stream is always directed to the primary of a primary-backup group, which is transparent to both the call processing client application and the call process (e.g., CP1, CP2, CP3, TIL1, TIL2, IS1, IS2).

AS noted above, the call processes on the call application nodes (CANs) are organized into a load sharing group. Each call process (e.g., CP1, CP2, CP3, TIL1, TIL2, IS1, IS2) is itself a primary-backup group. Both members of the primary-backup group may provide the service but only the primary of the group receives messages and thus actually provides the service. When a member of the group is selected as the primary, it registers one or more interface streams for the group. Each stream is a separate interface for some call processing service.

The call processing client application (e.g., CPC APP1, CPC

APP2) in MSC 140 receives a new call indication and uses the group service to select an interface with a call application node (i.e., server) to handle the new call. The call process on each server (CAN) is a member of a load sharing group and a particular call application node (CAN) is selected using a round-robin algorithm from the perspective of the call process client application. For the particular primary-backup group that is selected a session is returned to the call processing client application. When the session is established with the primary-backup call process server group, the call processing client application then opens an interface to a particular member (representing an interface to a primary-backup group) and obtains a session interface. Each call processing server sends a message related to the new call over the session interface. Any subsequent transactions associated with the call are sent over the same session object.

The call process server (i.e., primary-backup group) may send asynchronously messages over the session using one or more of the defined stream interfaces. The primary member of the call processing server group receives the transactions. The backup group member does not receive transactions. The primary group member sends updates to the backup group member. The primary group member decides when updates are sent to the backup group member.

The primary starts sending updates when a call has been answered. Prior to the call being answered, the call is defined as being a transient call. After the call has been answered, the call is defined as being a stable call.

5 If the primary group member should fail, then the backup group member becomes the new primary member. All transient call information during the fail-over period (the time between when the primary fails and the backup is changed to be the new primary) can be lost. All stable call information must be maintained by the  
10 backup. However, some stable call information may be lost if the backup has not received updates.

Advantageously, the present invention has no limitations on the scalability of the system and the system size is hidden from both and the primary-backup group server applications and call  
15 process client applications. The present invention eliminates any single point of failure in the system. Any failure within the system will not affect the system availability and performance.

New call application nodes (CANs) and additional primary-backup group server applications (e.g., CP1, CP2, CP3, TIL1, TIL2,  
20 IS1, IS2) may be added dynamically to the load sharing groups and can start servicing new call traffic. Call process client applications are not affected by the additions of new servers. If

a server should fail, its backup assumes responsibility for the load. This provides high availability for the servicing of each call and minimizes dropped calls.

FIGURE 4A is an illustration of server side internal group policy classes according to an exemplary embodiment of the present invention. FIGURE 4B is an illustration of a client side internal client policy architecture according to an exemplary embodiment of the present invention. FIGURE 4C is an illustration of a load sharing client side policy internal architecture according to an exemplary embodiment of the present invention.

Server applications create a primary-backup group and then join the primary-backup group. This action creates the server side policy containing the objects shown in FIGURE 4A. The group policy distributes invocations from clients, participates in a distributed election of the primary in the group, maintains group membership, and monitors for group member failures. Server applications join a load sharing group using a group adaptor object as a proxy member of the load sharing group. The group adaptor object is set with the name of the primary-backup group prior to joining the load sharing group.

Client applications establish a client interface to the load sharing group and begin by opening a session. The act of opening

a session utilizes a client side load sharing policy to select one of the members of the load sharing group. The internal load sharing client policy architecture is shown in FIGURE 4C. The session object itself encapsulates a client side policy that connects to a particular primary-backup group. The internal architecture for this client policy is shown in FIGURE 4B.

Call processing client application communicate with the selected server (which is the primary within a primary-backup group). As the primary call process receives messages from the call processing client application, the primary call process sends state updates to the corresponding backup call process. If the primary call process should fail, the backup call process is automatically selected as the new primary. During the fail-over period to the new primary, the call processing client application receives an indication of the failure and may retry the send until the new primary call process is ready to receive messages. This minimizes the lost message traffic during the fail-over period. Once the call processing client application is through with the session, the call processing client application may release the session.

The PBUNIGroupPolicy group policy has the following internal member:

1) PBUNIConfiguration - identifies the group policy name as being "PBUNI" and specifies the QoS requirements for the communication stack for this policy.

2) PolicyGroupMembership - maintains the membership for the group and provides a membership protocol for adding new members, removing members that have left, and marking members that have failed as "FAILED".

3) PBPolicyEventNotification - provides the behavior for event notifications, such as i) when a member joins the group (recovered), ii) leaves the group (left), iii) fails (failed), or iv) has a state change.

4) PBMemberStateControl - has the state machine for primary selection in the presence of joins, exits, and failures of group members. Each local instance of this class decides which member is the primary. It is possible, due to network partitions, that there can be more than one primary at the same time.

5) PBSessionControl - controls the session establishment between call processing client applications for a primary-backup group and the group members.

6) PBPolicyIOControl - provides the primary-backup policy behavior for multicasting and sending to group members.

7) GroupSendProtocol - provides the group member protocol



for sending to other members of the group and to clients of the group.

8) UnicastGroupInterface - is a group interface that provides separate interfaces to each capsule in which a group member resides.

PBUNIClientPolicy is a primary-backup client policy in which unicast links are used to communicate with the group. General sending is sent only to the primary member and is not redundantly sent to the backup member(s). The PBUNIClientPolicy has the following members:

1) ClientMembershipView - provides a local view of the group membership but unlike GroupMembershipView, does not participate in the protocol associated with group membership.

2) PBUNIClientPolicyIO - handles I/O over unicast links to the primary member.

3) GroupSendProtocol - provides the Group Member protocol for sending to other members of the group and to clients of the group.

4) ClientSessionControl - manages sessions on the client side with group members.

5) PBUNIClientStateControl - maintains a local view of which member is the primary in a primary-backup group.

6) ClientSessionControl - manages sessions on the client side with group members.

7) UnicastGroupInterface - provides separate interfaces to each capsule in which a group member resides.

5 LSIClientPolicy is a load sharing policy for group members who are themselves groups. LSIClientPolicy provides client-side round-robin selection of members when a session to a member is opened. Each session provides a group interface to a particular group. LSIClientPolicy also provides support of message throttling to each  
10 session. Throttling can vary from 0% (no messages) to 100% (all messages are sent that would normally be selected using round-robin scheduling). LSIClientPolicy overrides what is in the base ClientPolicy. LSIClientPolicy contains the following members:

1) LSIClientNotification - notifies both the ClientPolicy  
15 notifier and the local notifier of events.

2) LSIClientSessionControl - returns a session using a round-robin algorithm. The session provided is an interface to another group. LSISessionControl has a running index that is used to select a new session for each open session request.  
20 LSISessionControl has a list of known interfaces called "Member Known". Member Known is a map that is indexed by the Member ID and contains a SessionCount object which contains the actual session

and a reference count of the number of users of the session instance. The sessions in known are maintained even though members may leave the group. When members leave the group they are removed from being available but kept in known. This permits clients to  
5 continue to use the group interface even though they have left the load sharing group.

3) GroupSendProtocol - provides the Group Member protocol for sending to other members of the group and to clients of the group.

10 4) GroupInterface - is the interface class to the multicast and/or unicast protocol stack(s) that are utilized by the group interfaces.

5) ClientPolicyIO - is responsible for handling client I/O.

15 6) ClientStateControl - is used to control the event state of the group and to retrieve the event state of the group.

According to one embodiment of the present invention, MSC 140 utilizes the group service approach to provide a distributed set of identity servers in MSC 140 and wireless network 100. As stated above, IS1 and IS2 are identity server (IS) applications that  
20 provide services that are invoked by various client applications, such as CPC APP1 - CPC APPn. The IS1 and IS2 server applications are organized into primary-backup groups that are scalable

according to the number of processes and/or computing nodes (CANs) that are used. The IS1 and IS2 server applications are members of a load sharing group, namely LSG3. Call process server applications CP1-CP3 are clients with respect to LSG3 and the IS1 and IS2 server applications. Each call identity server application (e.g., IS1, IS2) is part of a primary-backup group server application.

Generally speaking, the present invention uses a set of N call identity servers (there is no limitation) that manage a range of call identities within the total range of call identities. Typically, a call identity is a 14 digit number. Thus, for example, for a set of ten servers, each call identity server application (e.g., IS1, IS2) manages a tenth of the total call identity range. When a new call is originated, a client application for the LSG3 identity server load sharing group requests an identity. The Client Policy for the identity server load sharing group then selects the least loaded server application according to one or more of the following criterion: 1) using a client side round-robin selection of one of the total list of available servers; or 2) selecting a server application that is least loaded, such as lower CPU utilization and/or lower memory utilization.

Once such a server is selected, the server (e.g., LS1) allocates one of the call identity numbers in the range it manages. The call identity number is thereafter used to identify the call. When the call ends, the call identity number is released by  
5 locating the proper server from the call identity itself, as each server manages a fixed range of numbers. The server application then makes the number available by marking it as unallocated.

The identity server application is related to the call distribution server application in that when the client application  
10 in MPU 210 selects a call processing application, the first thing that the call processing application does is act as a client to the identity server group to allocate a call identity number.

According to an advantageous embodiment of the present invention, there is a configurable range of available call identity  
15 values. When the identity server load sharing group (i.e., LSG3) is started, either IS1 or IS2 is elected "leader" and it assigns subsets of the call identity values to each of the available servers. Each of the IS1 and IS2 group server applications manages a contiguous range of call identity values.

20 When a client application (e.g., CP1, CP2, CP3) accesses the identity server load sharing group, the client application client policy is given a Resource Reference Table. The Resource Reference

Table is a look-up table organized according to a range of call identity values that are associated with an interface to the specific server. The table contents are automatically updated by the LSG3 identity server load sharing group as server configurations change.

When a call begins, a call identity is allocated at the client interface. The LSG3 client policy receives loading factors from the IS1 and IS2 servers. When a new call identity is allocated, LSG3 client policy takes into account the loading of the servers and selects the server that is least loaded. Any client application in the network may now read from or write to the associated data as long as each client application has the call identity value. The client application just specifies the call identity and the interface to the IS1 or IS2 server is returned. The client application can then read or write to the associated data or release the call identity value.

Although the present invention has been described in detail, those skilled in the art should understand that they may make various changes, substitutions and alterations herein without departing from the spirit and scope of the invention in its broadest form.